PORTAL USPTO

Subscribe (Full Service) Register (Limited Service, Free) Login

Search: The ACM Digital Library The Guide

"partition database", file context"

THE ACM DIGITAL LIBRARY

Feedback Report a problem Satisfaction survey

Terms used **partition database file context** Found **19,633** of **169,166**

Sort results by: relevance

Display results: expanded form

Save results to a Binder

Search Tips

Open results in a new window

Try an Advanced Search

Try this search in The ACM Guide

Results 1 - 20 of 200 Result page: **1** 2 3 4 5 6 7 8 9 10 next

Best 200 shown Relevance scale

1 An application of a context-aware file system

Christopher K. Hess, Roy H. Campbell

December 2003 **Personal and Ubiquitous Computing**, Volume 7 Issue 6

Publisher: Springer-Verlag

Full text available: pdf(383.26 KB) Additional Information: full citation, abstract, citings, index terms

Ubiquitous computing environments stretch the requirements of traditional infrastructures used to facilitate the development of applications. Activities are often supported by collections of applications, some of which are automatically launched with little or no human intervention. This task-driven environment challenges existing application construction and data management techniques. In this paper, we describe a file system that organises application data based on contextual information, impo ...

Keywords: Context, Data management, File systems, Operating systems, Ubiquitous computing spaces

2 Connections: using context to enhance file search

Craig A. N. Soules, Gregory R. Ganger

October 2005 **ACM SIGOPS Operating Systems Review , Proceedings of the twentieth ACM symposium on Operating systems principles SOSP '05**, Volume 39 Issue 5

Publisher: ACM Press

Full text available: pdf(360.72 KB) Additional Information: full citation, abstract, references, index terms

Connections is a file system search tool that combines traditional content-based search with context information gathered from user activity. By tracing file system calls, Connections can identify temporal relationships between files and use them to expand and reorder traditional content search results. Doing so improves both recall (reducing false-positives) and precision (reducing false-negatives). For example, Connections improves the average recall (from 13% to 22%) and precision (from 23% t ...

Keywords: context, file system search, successor models

3 Piranha: a scalable architecture based on single-chip multiprocessing

Luiz André Barroso, Kourosh Gharachorloo, Robert McNamara, Andreas Nowatzyk, Shaz Qadeer, Barton Sano, Scott Smith, Robert Stets, Ben Verghese

May 2000 **ACM SIGARCH Computer Architecture News , Proceedings of the 27th annual international symposium on Computer architecture ISCA '00**, Volume 28 Issue 2
Publisher: ACM Press

Full text available: pdf(191.10 KB) Additional Information: full citation, abstract, references, citings, index terms

The microprocessor industry is currently struggling with higher development costs and longer design times that arise from exceedingly complex processors that are pushing the limits of instruction-level parallelism. Meanwhile, such designs are especially ill suited for important commercial applications, such as on-line transaction processing (OLTP), which suffer from large memory stall times and exhibit little instruction-level parallelism. Given that commercial applications constitute by fa ...

4 Memory system characterization of commercial workloads

Luiz André Barroso, Kourosh Gharachorloo, Edouard Bugnion
April 1998 **ACM SIGARCH Computer Architecture News , Proceedings of the 25th annual international symposium on Computer architecture ISCA '98**, Volume 26 Issue 3
Publisher: IEEE Computer Society, ACM Press

Full text available: pdf(1.68 MB) Publisher Site Additional Information: full citation, abstract, references, citings, index terms

Commercial applications such as databases and Web servers constitute the largest and fastest-growing segment of the market for multiprocessor servers. Ongoing innovations in disk subsystems, along with the ever increasing gap between processor and memory speeds, have elevated memory system design as the critical performance factor for such workloads. However, most current server designs have been optimized to perform well on scientific and engineering workloads, potentially leading to design dec ...

5 Ubiquitous computing (UC): Intelligent file management in ubiquitous environments

Kartik Vishwanath, Arvind Gautam, Yugyung Lee
March 2005 **Proceedings of the 2005 ACM symposium on Applied computing**
Publisher: ACM Press

Full text available: pdf(196.84 KB) Additional Information: full citation, abstract, references, index terms

The paradigm of Ubiquitous computing seeks to build a computing environment that responds to user context. An ideal file system for the Ubiquitous environment is one that can successfully recognize the present context and automate file management. The intelligence in the Ubiquitous file management is achieved by applying a heuristics based clustering approach to the system. The applied heuristics are those that are used on file attributes by users to manually manage files in a traditional file s ...

Keywords: classification, context, file systems, ubiquitous computing

6 Using event contexts and matching constraints to monitor software processes

Naser S. Barghouti, Balachander Krishnamurthy
April 1995 **Proceedings of the 17th international conference on Software engineering**
Publisher: ACM Press

Full text available: pdf(941.11 KB) Additional Information: full citation, references, citings, index terms

7 Dynamic construction of animated help from application context

Piyawadee Sukaviriya
January 1988 **Proceedings of the 1st annual ACM SIGGRAPH symposium on User**

Interface Software
Publisher: ACM Press

Full text available: pdf(1.64 MB) Additional Information: full citation, abstract, references, citings, index terms

Help provided as traditional text descriptions has become incompatible with graphical interfaces. Animation suggests a better association between help and a graphical interface. This paper describes a prototype system implemented to demonstrate the use of dynamic scenarios as help. A scenario animates the execution of a task as a sequence of steps in the actual interface and work context. Each scenario is dynamically generated depending on the current work context of the user. The system re ...

8 Register relocation: flexible contexts for multithreading

Carl A. Waldspurger, William E. Weihl
May 1993 **ACM SIGARCH Computer Architecture News , Proceedings of the 20th annual international symposium on Computer architecture ISCA '93**, Volume 21 Issue 2
Publisher: ACM Press

Full text available: pdf(1.06 MB) Additional Information: full citation, abstract, references, citings, index terms

Multithreading is an important technique that improves processor utilization by allowing computation to be overlapped with the long latency operations that commonly occur in multiprocessor systems. This paper presents register relocation, a new mechanism that efficiently supports flexible partitioning of the register file into variable-size contexts with minimal hardware support. Since the number of registers required by thread contexts varies, this flexibility permits a better utilization ...

9 Optimizing throughout in a workstation-based network file system over a high bandwidth local area network

Theodore Faber
January 1998 **ACM SIGOPS Operating Systems Review**, Volume 32 Issue 1
Publisher: ACM Press

Full text available: pdf(858.90 KB) Additional Information: full citation, abstract, index terms

This paper describes methods of optimizing a client/server network file system to advantage of high bandwidth local area networks in a conventional distributed computing environment. The environment contains hardware that removes network and disk bandwidth bottlenecks. The remaining bottlenecks at clients include excessive context switching, inefficient data translation, and cumbersome data encaps ulation methods. When these are removed, the null-write performance of a current implementation of S ...

10 Full Papers: Exposing document context in the personal web

David Wolber, Michael Kepe, Igor Ranitovic
January 2002 **Proceedings of the 7th international conference on Intelligent user interfaces**
Publisher: ACM Press

Full text available: pdf(295.10 KB) Additional Information: full citation, abstract, references, citings, index terms

Reconnaissance agents show context by displaying documents with similar content to the one(s) the user currently has open. *Research paper search engines* show context by displaying documents that cite or are cited by the currently open document(s). We present a tool that applies such ideas to the *personal web*, that is, the space rooted in user documents but tightly connected to web documents as well. The tool organizes the personal web with a single topic hierarchy based on d ...

Keywords: context, information navigation, personal web, recommender, reconnaissance

11 Threads: Balancing register pressure and context-switching delays in ASTI systems

Siddhartha Shivshankar, Sunil Vangara, Alexander G. Dean

September 2005 **Proceedings of the 2005 international conference on Compilers, architectures and synthesis for embedded systems CASES '05**

Publisher: ACM Press

Full text available: pdf(261.60 KB) Additional Information: full citation, abstract, references, index terms

This paper makes two contributions to Asynchronous Software Thread Integration (ASTI). First, it presents methods to calculate worst-case secondary thread performance statically. This will enable real-time performance guarantees for the system in future work. Second, it improves the run-time performance of integrated threads by partitioning the register file, allowing faster coroutine calls. Determining the ideal partitioning of the register file is non-trivial if the registers are heterogeneous ...

Keywords: asynchronous software thread integration, fine-grain concurrency, hardware to software migration, software-implemented-communication protocols

12 CUROCO: a distributed architecture for the dynamic generation, composition and use of context in highly dynamic and heterogeneous environments [extended version]

Gorka Guardiola Muzquiz

October 2004 **Proceedings of the 1st international doctoral symposium on Middleware**

Publisher: ACM Press

Full text available: pdf(67.20 KB) Additional Information: full citation, abstract, references, index terms

The use of context is necessary for a rich interaction with the users and adaptation of the system to their needs. However, how to integrate context with the applications and the system is yet an unsolved problem. Looking at other services provided by our computing platforms, we ask ourselves: How can context be made a system service?. Making the context an application neutral system service would permit to share and update it from different sources an applications, making the user able to in ...

Keywords: CUROCO, context aware, pervasive, sentient computing, ubiquitous

13 Measurement and analysis of locality phases in file referencing behaviour

Shikharesh Majumdar, Richard B. Bunt

May 1986 **ACM SIGMETRICS Performance Evaluation Review , Proceedings of the 1986 ACM SIGMETRICS joint international conference on Computer performance modelling, measurement and evaluation SIGMETRICS '86/PERFORMANCE '86**, Volume 14 Issue 1

Publisher: ACM Press

Full text available: pdf(1.39 MB) Additional Information: full citation, abstract, references, citings, index terms

Recent research has demonstrated the existence of locality in short-term file referencing behaviour. A detailed study of the dynamic characteristics of file referencing is presented in this paper. The concept of Bounded Locality Intervals from the field of program behaviour has been used to model the locality phases of file referencing behaviour. The model is found to be powerful both from a descriptive point of view and from the perspective of understanding the performance implicat ...

14 Distributed file systems: concepts and examples

Eliezer Levy, Abraham Silberschatz

December 1990 **ACM Computing Surveys (CSUR)**, Volume 22 Issue 4
Publisher: ACM Press
Full text available: pdf(5.33 MB) Additional Information: full citation, abstract, references, citings, index terms, review

The purpose of a distributed file system (DFS) is to allow users of physically distributed computers to share data and storage resources by using a common file system. A typical configuration for a DFS is a collection of workstations and mainframes connected by a local area network (LAN). A DFS is implemented as part of the operating system of each of the connected computers. This paper establishes a viewpoint that emphasizes the dispersed structure and decentralization of both data and con ...

15 LegionFS: a secure and scalable file system supporting cross-domain high-performance applications
Brian S. White, Michael Walker, Marty Humphrey, Andrew S. Grimshaw
November 2001 **Proceedings of the 2001 ACM/IEEE conference on Supercomputing (CDROM)**
Publisher: ACM Press
Full text available: pdf(499.88 KB) Additional Information: full citation, abstract, references, citings, index terms

Realizing that current file systems can not cope with the diverse requirements of wide-area collaborations, researchers have developed data access facilities to meet their needs. Recent work has focused on comprehensive data access architectures. In order to fulfill the evolving requirements in this environment, we suggest a more fully-integrated architecture built upon the fundamental tenets of naming, security, scalability, extensibility, and adaptability. These form the underpinning of the Le ...

16 A reduced register file for RISC architectures
Miquel Huguet, Tomás Lang
September 1985 **ACM SIGARCH Computer Architecture News**, Volume 13 Issue 4
Publisher: ACM Press
Full text available: pdf(800.03 KB) Additional Information: full citation, citings, index terms

17 Special feature on MOBICOM 2004 posters: A context based storage system for mobile computing applications
Sharat Khungar, Jukka Riekki
January 2005 **ACM SIGMOBILE Mobile Computing and Communications Review**, Volume 9 Issue 1
Publisher: ACM Press
Full text available: pdf(201.80 KB) Additional Information: full citation, abstract, references, index terms

In this paper, we describe a novel context based storage system that use context to manage user data and make it available to him based on his situation. First, we examine several existing systems that use context with documents. Subsequently, a new storage system is presented that uses context to aid in the capture of and access to documents in mobile environment. We describe file browser and calendar applications that we have developed within our mobile computing infrastructure utilizing the f ...

18 Extensible file systems in spring
Yousef A. Khalidi, Michael N. Nelson
December 1993 **ACM SIGOPS Operating Systems Review , Proceedings of the fourteenth ACM symposium on Operating systems principles SOSP '93**, Volume 27 Issue 5
Publisher: ACM Press

Full text available: pdf(1.47 MB) Additional Information: full citation, abstract, references, citings, index terms

In this paper we describe an architecture for extensible file systems. The architecture enables the extension of file system functionality by composing (or stacking) new file systems on top of existing file systems. A file system that is stacked on top of an existing file system can access the existing file system's files via a well-defined naming interface and can share the same underlying file data in a coherent manner. We describe extending file systems in the context of the Spring operating ...

19 A logical view of structured files

Serge Abiteboul, Sophie Cluet, Tova Milo

May 1998 **The VLDB Journal — The International Journal on Very Large Data Bases**, Volume 7 Issue 2

Publisher: Springer-Verlag New York, Inc.

Full text available: pdf(288.29 KB) Additional Information: full citation, abstract, citings, index terms

Structured data stored in files can benefit from standard database technology. In particular, we show here how such data can be queried and updated using declarative database languages. We introduce the notion of *structuring schema*, which consists of a grammar annotated with database programs. Based on a structuring schema, a file can be viewed as a database structure, queried and updated as such. For *queries*, we show that almost standard database optimization techniques can be use ...

Keywords: Database, File system, Query, Query and update optimization, Textual data, Update

20 Posters: A context based storage for ubiquitous computing applications

Sharat Khungar, Jukka Riekki

November 2004 **Proceedings of the 2nd European Union symposium on Ambient intelligence EUSAI '04**

Publisher: ACM Press

Full text available: pdf(444.96 KB) Additional Information: full citation, abstract, references

Context-aware systems are computing systems that provide relevant services and information to users based on their situational conditions [1]. Here, first we present a context based storage system that uses context to aid in the capture and access of documents in ubiquitous environment. Then we describe a set of context-aware applications that we have developed within our ubiquitous computing infrastructure utilizing the features of context based storage. Novel features of our system include ...

Keywords: context-aware, data management, ubiquitous computing

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2006 ACM, Inc.
Terms of Usage Privacy Policy Code of Ethics Contact Us

Useful downloads: Adobe Acrobat QuickTime Windows Media Player Real Player

PORTAL USPTO

Subscribe (Full Service) Register (Limited Service, Free) Login

Search: ◉ The ACM Digital Library ○ The Guide

"partition database", file" "lowest"

THE ACM DIGITAL LIBRARY

Feedback Report a problem Satisfaction survey

Terms used **partition database file lowest** Found **5,660** of **169,166**

Sort results by: relevance

Display results: expanded form

Save results to a Binder

Search Tips

☐ Open results in a new window

Try an Advanced Search

Try this search in The ACM Guide

Results 1 - 20 of 200 Result page: **1** 2 3 4 5 6 7 8 9 10 next

Best 200 shown Relevance scale

1 Memory system characterization of commercial workloads

Luiz André Barroso, Kourosh Gharachorloo, Edouard Bugnion

April 1998 **ACM SIGARCH Computer Architecture News , Proceedings of the 25th annual international symposium on Computer architecture ISCA '98**, Volume 26 Issue 3

Publisher: IEEE Computer Society, ACM Press

Full text available: pdf(1.68 MB) Publisher Site Additional Information: full citation, abstract, references, citings, index terms

Commercial applications such as databases and Web servers constitute the largest and fastest-growing segment of the market for multiprocessor servers. Ongoing innovations in disk subsystems, along with the ever increasing gap between processor and memory speeds, have elevated memory system design as the critical performance factor for such workloads. However, most current server designs have been optimized to perform well on scientific and engineering workloads, potentially leading to design dec ...

2 Multiple-banked register file architectures

José-Lorenzo Cruz, Antonio González, Mateo Valero, Nigel P. Topham

May 2000 **ACM SIGARCH Computer Architecture News , Proceedings of the 27th annual international symposium on Computer architecture ISCA '00**, Volume 28 Issue 2

Publisher: ACM Press

Full text available: pdf(106.23 KB) Additional Information: full citation, abstract, references, citings, index terms

The register file access time is one of the critical delays in current superscalar processors. Its impact on processor performance is likely to increase in future processor generations, as they are expected to increase the issue width (which implies more register ports) and the size of the instruction window (which implies more registers), and to use some kind of multithreading. Under this scenario, the register file access time could be a dominant delay and a pipelined implementation would ...

Keywords: bypass logic, dynamically-scheduled processor, register file architecture, register file cache

3 Sorting nonredundant files—techniques used in the FACT compiler

John B. Glore

May 1963 **Communications of the ACM**, Volume 6 Issue 5

Publisher: ACM Press

Full text available: pdf(1.02 MB) Additional Information: full citation, abstract, citings

Some typical file structures, including some called "non-redundant," are examined, and the methods used in FACT to sort such files are discussed.

4 A file organization for cluster-based retrieval

W. Bruce Croft

May 1978 **ACM SIGIR Forum , Proceedings of the 1st annual international ACM SIGIR conference on Information storage and retrieval SIGIR '78**, Volume 13 Issue 1

Publisher: ACM Press

Full text available: pdf(659.26 KB) Additional Information: full citation, abstract, references, index terms

A file organization for cluster-based retrieval is presented and tested. This file organization is based on the bottom-up search which, in contrast to the more usual top-down search, starts at the lowest level of a cluster hierarchy (the documents) and looks at progressively larger clusters. This approach enables most of the efficiency problems previously associated with clustered file organizations to be avoided. There are two parts to this file organization - a compact cluster hierarchy r ...

5 The BANG file: A new kind of grid file

Michael Freeston

December 1987 **ACM SIGMOD Record , Proceedings of the 1987 ACM SIGMOD international conference on Management of data SIGMOD '87**, Volume 16 Issue 3

Publisher: ACM Press

Full text available: pdf(867.07 KB) Additional Information: full citation, abstract, references, citings, index terms

A new multi-dimensional file structure has been developed in the course of a project to devise ways of improving the support for interactive queries to database and knowledge bases. Christened the 'BANG' file - a Balanced And Nested Grid - the new structure is of the 'grid file' type, but is fundamentally different from previous grid file designs in that it does not share their common underlying properties. It has a tree-structured directory which has the self-balancing property of a B-tree ...

6 Mobile services: DeltaCast: efficient file reconciliation in wireless broadcast systems

Julian Chesterfield, Pablo Rodriguez

June 2005 **Proceedings of the 3rd international conference on Mobile systems, applications, and services MobiSys '05**

Publisher: ACM Press

Full text available: pdf(214.15 KB) Additional Information: full citation, abstract, references

Recently, there has been an increasing interest in wireless broadcast systems as a means to enable scalable content delivery to large numbers of mobile users. However, gracefully providing efficient reconciliation of different versions of a file over such broadcast channels still remains a challenge. Such systems often lack a feedback channel and consequently updates cannot be easily tailored to a specific user. Moreover, given the potentially large number of possible versions of a file, it is i ...

7 File organizations and access methods for CLV disks

S. Christodoulakis, D. A. Ford

December 1988 **ACM SIGIR Forum , Proceedings of the 12th annual international ACM SIGIR conference on Research and development in information retrieval SIGIR '89**, Volume 23 Issue 1-2

Publisher: ACM Press

Full text available: pdf(1.08 MB) Additional Information: full citation, abstract, references, citings, index terms

A large and important class of optical disc technology are CLV format discs such as CD ROM and WORM. In this paper, we examine the issues related to the implementation and performance of several different file organizations on CLV format optical discs such as CD ROM and WORM. The organizations examined are based on hashing and trees. The CLV recording scheme is shown to be a good environment for efficiently implementing hashing. Single seek access and storage utilization levels a ...

8 Design of a multi-level file management system

Edward W. Ver Hoef

January 1966 **Proceedings of the 1966 21st national conference**

Publisher: ACM Press

Full text available: pdf(1.08 MB) Additional Information: full citation, abstract, references, citings, index terms

This paper describes the file handling system developed and being implemented as a part of INTIPS, (INTe- grated Information Processing System) under the aegis of Rome Air Development Center. This file system addresses itself to the physical problem of the storage and retrieval of fixed length blocks and their organization into higher-order structures. It does not concern itself with the contents of these blocks. (For an elaboration of this distinction, see reference 1.)

9 Generation and search of clustered files

G. Salton, A. Wong

December 1978 **ACM Transactions on Database Systems (TODS)**, Volume 3 Issue 4

Publisher: ACM Press

Full text available: pdf(1.78 MB) Additional Information: full citation, abstract, references, citings, index terms

A classified, or clustered file is one where related, or similar records are grouped into classes, or clusters of items in such a way that all items within a cluster are jointly retrievable. Clustered files are easily adapted to broad and narrow search strategies, and simple file updating methods are available. An inexpensive file clustering method applicable to large files is given together with appropriate file search methods. An abstract model is then introduced to predict the retrieval ...

Keywords: automatic classification, cluster searching, clustered files, fast classification, file organization, probabilistic models

10 Implementation of the GIST Geographic Base File

Robert Amsterdam

January 1971 **Proceedings of the 1971 26th annual conference**

Publisher: ACM Press

Full text available: pdf(602.43 KB) Additional Information: full citation, abstract, references, index terms

GIST, New York City's Geographic Information System, is designed to unify many items of basic information describing the physical, social and environmental features of New York City. Files and programs are being developed for use by all agencies of the city government. The principal files are the Geographic Base File and the Building/Lot File. The GIST Geographic Base File (GBF) has been implemented ...

Keywords: Address matching, Census data, Geographic base file, Mapping by computer, Municipal information

11 Organization of clustered files for consecutive retrieval

J S. Deogun, V V. Raghavan, T K.W. Tsou

December 1984 **ACM Transactions on Database Systems (TODS)**, Volume 9 Issue 4

Publisher: ACM Press

Full text available: pdf(1.79 MB) Additional Information: full citation, abstract, references, index terms, review

This paper studies the problem of storing single-level and multilevel clustered files. Necessary and sufficient conditions for a single-level clustered file to have the consecutive retrieval property (CRP) are developed. A linear time algorithm to test the CRP for a given clustered file and to identify the proper arrangement of objects, if CRP exists, is presented. For the single-level clustered files that do not have CRP, it is shown that the problem of identifying a storage organization w ...

12 Measurement and analysis of locality phases in file referencing behaviour

Shikharesh Majumdar, Richard B. Bunt

May 1986 **ACM SIGMETRICS Performance Evaluation Review , Proceedings of the 1986 ACM SIGMETRICS joint international conference on Computer performance modelling, measurement and evaluation SIGMETRICS '86/PERFORMANCE '86**, Volume 14 Issue 1

Publisher: ACM Press

Full text available: pdf(1.39 MB) Additional Information: full citation, abstract, references, citings, index terms

Recent research has demonstrated the existence of locality in short-term file referencing behaviour. A detailed study of the dynamic characteristics of file referencing is presented in this paper. The concept of Bounded Locality Intervals from the field of program behaviour has been used to model the locality phases of file referencing behaviour. The model is found to be powerful both from a descriptive point of view and from the perspective of understanding the performance implicat ...

13 Denial-of-service resilience in peer-to-peer file sharing systems

D. Dumitriu, E. Knightly, A. Kuzmanovic, I. Stoica, W. Zwaenepoel

June 2005 **ACM SIGMETRICS Performance Evaluation Review , Proceedings of the 2005 ACM SIGMETRICS international conference on Measurement and modeling of computer systems SIGMETRICS '05**, Volume 33 Issue 1

Publisher: ACM Press

Full text available: pdf(245.14 KB) Additional Information: full citation, abstract, references, citings, index terms

Peer-to-peer (p2p) file sharing systems are characterized by highly replicated content distributed among nodes with enormous aggregate resources for storage and communication. These properties alone are not sufficient, however, to render p2p networks immune to denial-of-service (DoS) attack. In this paper, we study, by means of analytical modeling and simulation, the resilience of p2p file sharing systems against DoS attacks, in which malicious nodes respond to queries with erroneous responses. ...

Keywords: denial of service, file pollution, network-targeted attacks, peer-to-peer

14 Dynamic maintenance of data distribution for selectivity estimation

Kyu Young Whang, Sang Wook Kim, Gio Wiederhold

January 1994 **The VLDB Journal — The International Journal on Very Large Data Bases**, Volume 3 Issue 1

Publisher: Springer-Verlag New York, Inc.

Full text available: pdf(1.09 MB) Additional Information: full citation, abstract, references, citings

We propose a new dynamic method for multidimensional selectivity estimation for range queries that works accurately independent of data distribution. Good estimation of selectivity is important for query optimization and physical database design. Our method employs the multilevel grid file (MLGF) for accurate estimation of multidimensional data distribution. The MLGF is a dynamic, hierarchical, balanced, multidimensional file structure that gracefully adapts to nonuniform and correlated distribu ...

Keywords: multidimensional file structure, multilevel grid files, physical database design, query optimization

15 Comparison of loading costs for various file organizations

D. E. Strickland, J. D. Powell

April 1978 **Proceedings of the 16th annual Southeast regional conference**

Publisher: ACM Press

Full text available: pdf(327.00 KB) Additional Information: full citation, abstract, references

An algorithm is presented which determines the loading costs for the following file organizations: Sequential, Indexed Sequential, Direct, Multilist, and Inverted file structures. S. B. Yao has developed a hierarchical access model which provides a single model incorporating the above organizations. This model has been used for a dynamic reorganization algorithm, and performance evaluation of data base structures. This paper uses Yao's model as a basis for developing an algorithm to determine th ...

16 Feasibility of a serverless distributed file system deployed on an existing set of desktop PCs

William J. Bolosky, John R. Douceur, David Ely, Marvin Theimer

June 2000 **ACM SIGMETRICS Performance Evaluation Review , Proceedings of the 2000 ACM SIGMETRICS international conference on Measurement and modeling of computer systems SIGMETRICS '00**, Volume 28 Issue 1

Publisher: ACM Press

Full text available: pdf(946.00 KB) Additional Information: full citation, abstract, references, citings, index terms

We consider an architecture for a serverless distributed file system that does not assume mutual trust among the client computers. The system provides security, availability, and reliability by distributing multiple encrypted replicas of each file among the client machines. To assess the feasibility of deploying this system on an existing desktop infrastructure, we measure and analyze a large set of client machines in a commercial environment. In particular, we measure and report results on ...

Keywords: analytical modeling, availability, feasibility analysis, personal computer usage data, reliability, security, serverless distributed file system architecture, trust, workload characterization

17 A caching file system for a programmer's workstation

Michael D. Schroeder, David K. Gifford, Roger M. Needham

December 1985 **ACM SIGOPS Operating Systems Review , Proceedings of the tenth ACM symposium on Operating systems principles SOSP '85**, Volume 19 Issue 5

Publisher: ACM Press

Full text available: pdf(768.75 KB) Additional Information: full citation, references, citings, index terms

18 Disconnected operation in the Coda file system

James J. Kistler, M. Satyanarayanan
September 1991 **ACM SIGOPS Operating Systems Review , Proceedings of the thirteenth ACM symposium on Operating systems principles SOSP '91**, Volume 25 Issue 5
Publisher: ACM Press
Full text available: pdf(1.39 MB) Additional Information: full citation, abstract, references, citings, index terms

Disconnected operation is a mode of operation that enables a client to continue accessing critical data during temporary failures of a shared data repository. An important, though not exclusive, application of disconnected operation is in supporting portable computers. In this paper, we show that disconnected operation is feasible, efficient and usable by describing its design and implementation in the Coda File System. The central idea behind our work is that *caching of data*, now ...

19 Technical and social components of peer-to-peer computing: An end-user perspective on file-sharing systems
Jintae Lee
February 2003 **Communications of the ACM**, Volume 46 Issue 2
Publisher: ACM Press
Full text available: pdf(91.50 KB) html(25.44 KB) Additional Information: full citation, abstract, references, index terms

P2P file-sharing systems enable their users to share files directly among themselves without the need for a central file server. They form one of the most well-known categories of P2P systems, thanks largely to the Napster controversy and its appeal to the large potential user base. At its peak, Napster boasted a registered user base of 70 million [9] and 1.57 million simultaneous users. Now, after Napster's downfall, over 50 systems have taken its place. The files shared through these systems i ...

20 File organization and evaluation: Elements of the randomized combinatorial file structure
Richard A. Gustafson
April 1971 **Proceedings of the 1971 international ACM SIGIR conference on Information storage and retrieval**
Publisher: ACM Press
Full text available: pdf(730.19 KB) Additional Information: full citation, abstract, references, citings

A file structure designed to provide rapid, random access with minimum storage overhead is presented. Storage and retrieval are achieved by direct attribute combination-to-address transformation thereby negating the necessity for large file dictionaries or list-pointer structures. The attribute combination-to-address transformation is conceptually similar to key-to-address transformation techniques, but the transformation is not limited to operations on a single key but operates on the combi nati ...

Keywords: attribute combination-to-address transformation, combinatorial file, data structures, file access method, file organization, file structures, hash-coding, key-to-address transformation, multiple attribute retrieval, storage and retrieval system

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2006 ACM, Inc.
Terms of Usage Privacy Policy Code of Ethics Contact Us

Useful downloads: Adobe Acrobat QuickTime Windows Media Player Real Player

PORTAL USPTO
Subscribe (Full Service) Register (Limited Service, Free) Login
Search: The ACM Digital Library The Guide
partition or database or file lowest or smallest
THE ACM DIGITAL LIBRARY
Feedback Report a problem Satisfaction survey

Terms used **partition** or **database** or **file** **lowest** or **smallest** Found **35,301** of **169,166**

Sort results by: relevance
Display results: expanded form
Save results to a Binder
Search Tips
Open results in a new window

Try an Advanced Search
Try this search in The ACM Guide

Results 1 - 20 of 200 Result page: **1** 2 3 4 5 6 7 8 9 10 next
Best 200 shown Relevance scale

1 Searching in high-dimensional spaces: Index structures for improving the performance of multimedia databases
Christian Böhm, Stefan Berchtold, Daniel A. Keim
September 2001 **ACM Computing Surveys (CSUR)**, Volume 33 Issue 3
Publisher: ACM Press
Full text available: pdf(1.39 MB) Additional Information: full citation, abstract, references, citings, index terms

During the last decade, multimedia databases have become increasingly important in many application areas such as medicine, CAD, geography, and molecular biology. An important research issue in the field of multimedia databases is the content-based retrieval of similar multimedia objects such as images, text, and videos. However, in contrast to searching data in a relational database, a content-based retrieval requires the search of similar objects as a basic functionality of the database system ...

Keywords: Index structures, indexing high-dimensional data, multimedia databases, similarity search

2 Join processing in relational databases
Priti Mishra, Margaret H. Eich
March 1992 **ACM Computing Surveys (CSUR)**, Volume 24 Issue 1
Publisher: ACM Press
Full text available: pdf(4.42 MB) Additional Information: full citation, abstract, references, citings, index terms, review

The join operation is one of the fundamental relational database query operations. It facilitates the retrieval of information from two different relations based on a Cartesian product of the two relations. The join is one of the most diffidult operations to implement efficiently, as no predefined links between relations are required to exist (as they are with network and hierarchical systems). The join is the only relational algebra operation that allows the combining of related tuples fro ...

Keywords: database machines, distributed processing, join, parallel processing, relational algebra

3 SURVEY: Decoupling partitioning and grouping: Overcoming shortcomings of spatial

indexing with bucketing
Hanan Samet
December 2004 **ACM Transactions on Database Systems (TODS)**, Volume 29 Issue 4
Publisher: ACM Press
Full text available: pdf(446.42 KB) Additional Information: full citation, abstract, references, index terms

The principle of decoupling the partitioning and grouping processes that form the basis of most spatial indexing methods that use tree directories of buckets is explored. The decoupling is designed to overcome the following drawbacks of traditional solutions:(1) multiple postings in disjoint space decomposition methods that lead to balanced trees such as the hB-tree where a node split in the event of node overflow may be such that one of the children of the node that was split becomes a child of ...

Keywords: BV-trees, PK-trees, R-trees, Spatial indexing, decoupling, object hierarchies, space decomposition

4 Speeding up construction of PMR quadtree-based spatial indexes
Gisli R. Hjaltason, Hanan Samet
October 2002 **The VLDB Journal — The International Journal on Very Large Data Bases**, Volume 11 Issue 2
Publisher: Springer-Verlag New York, Inc.
Full text available: pdf(355.72 KB) Additional Information: full citation, abstract, citings, index terms

Spatial indexes, such as those based on the qu adtree, are important in spatial databases for efficient execution of queries involving spatial constraints, especially when the queries involve spatial joins. In this paper we present a number of techniques for speeding up the construction of quadtree-based spatial indexes, specifically the PMR quadtree, which can index arbitrary spatial data. We assume a quadtree implementation using the "linear quadtree", a disk-resident representation ...

Keywords: Bulk-loading, I/O, Spatial indexing

5 Parallelism and concurrency control performance in distributed database machines
Michael J. Carey, Miron Livny
June 1989 **ACM SIGMOD Record , Proceedings of the 1989 ACM SIGMOD international conference on Management of data SIGMOD '89**, Volume 18 Issue 2
Publisher: ACM Press
Full text available: pdf(1.56 MB) Additional Information: full citation, abstract, references, citings, index terms

While several distributed (or 'shared nothing') database machines exist in the form of prototypes or commercial products, and a number of distributed concurrency control algorithms are available, the effect of parallelism on concurrency control performance has received little attention. This paper examines the interplay between parallelism and transaction performance in a distributed database machine context. Four alternative concurrency control algorithms are considered, including two-phas ...

6 The BANG file: A new kind of grid file
Michael Freeston
December 1987 **ACM SIGMOD Record , Proceedings of the 1987 ACM SIGMOD international conference on Management of data SIGMOD '87**, Volume 16 Issue 3
Publisher: ACM Press
Full text available: pdf(867.07 KB) Additional Information: full citation, abstract, references, citings, index terms

A new multi-dimensional file structure has been developed in the course of a project to devise ways of improving the support for interactive queries to database and knowledge bases. Christened the 'BANG' file - a Balanced And Nested Grid - the new structure is of the 'grid file' type, but is fundamentally different from previous grid file designs in that it does not share their common underlying properties. It has a tree-structured directory which has the self-balancing property of a B-tree ...

7 Distributed file systems: concepts and examples

Eliezer Levy, Abraham Silberschatz
December 1990 **ACM Computing Surveys (CSUR)**, Volume 22 Issue 4
Publisher: ACM Press

Full text available: pdf(5.33 MB) Additional Information: full citation, abstract, references, citings, index terms, review

The purpose of a distributed file system (DFS) is to allow users of physically distributed computers to share data and storage resources by using a common file system. A typical configuration for a DFS is a collection of workstations and mainframes connected by a local area network (LAN). A DFS is implemented as part of the operating system of each of the connected computers. This paper establishes a viewpoint that emphasizes the dispersed structure and decentralization of both data and con ...

8 External memory algorithms and data structures: dealing with **massive data**

Jeffrey Scott Vitter
June 2001 **ACM Computing Surveys (CSUR)**, Volume 33 Issue 2
Publisher: ACM Press

Full text available: pdf(828.46 KB) Additional Information: full citation, abstract, references, citings, index terms

Data sets in large applications are often too massive to fit completely inside the computers internal memory. The resulting input/output communication (or I/O) between fast internal memory and slower external memory (such as disks) can be a major performance bottleneck. In this article we survey the state of the art in the design and analysis of external memory (or EM) algorithms and data structures, where the goal is to exploit locality in order to reduce the I/O costs. We consider a varie ...

Keywords: B-tree, I/O, batched, block, disk, dynamic, extendible hashing, external memory, hierarchical memory, multidimensional access methods, multilevel memory, online, out-of-core, secondary storage, sorting

9 Dynamic partitioning of signature files

P. Zezula, F. Rabitti, P. Tiberio
October 1991 **ACM Transactions on Information Systems (TOIS)**, Volume 9 Issue 4
Publisher: ACM Press

Full text available: pdf(2.22 MB) Additional Information: full citation, references, citings, index terms, review

10 Join algorithm costs revisited

Evan P. Harris, Kotagiri Ramamohanarao
January 1996 **The VLDB Journal — The International Journal on Very Large Data Bases**, Volume 5 Issue 1
Publisher: Springer-Verlag New York, Inc.

Full text available: pdf(329.00 KB) Additional Information: full citation, abstract, citings, index terms

A method of analysing join algorithms based upon the time required to access, transfer

and perform the relevant CPU-based operations on a disk page is proposed. The costs of variations of several of the standard join algorithms, including nested block, sort-merge, GRACE hash and hybrid hash, are presented. For a given total buffer size, the cost of these join algorithms depends on the parts of the buffer allocated for each purpose. For example, when joining two relations using the nested block j ...

Keywords: Join algorithms, Minimisation, Optimal buffer allocation

11 Multidimensional access methods

Volker Gaede, Oliver Günther
June 1998 **ACM Computing Surveys (CSUR)**, Volume 30 Issue 2
Publisher: ACM Press
Full text available: pdf(1.05 MB) Additional Information: full citation, abstract, references, citings, index terms

Search operations in databases require special support at the physical level. This is true for conventional databases as well as spatial databases, where typical search operations include the point query (find all objects that contain a given search point) and the region query (find all objects that overlap a given search region). More than ten years of spatial database research have resulted in a great variety of multidimensional access methods to support ...

Keywords: data structures, multidimensional access methods

12 Special issue: Game-playing programs: theory and practice

M. A. Bramer
April 1982 **ACM SIGART Bulletin**, Issue 80
Publisher: ACM Press
Full text available: pdf(9.23 MB) Additional Information: full citation, abstract

This collection of articles has been brought together to provide SIGART members with an overview of Artificial Intelligence approaches to constructi ng game-playing programs. Papers on both theory and practice are included.

13 Fast detection of communication patterns in distributed executions

Thomas Kunz, Michiel F. H. Seuren
November 1997 **Proceedings of the 1997 conference of the Centre for Advanced Studies on Collaborative research**
Publisher: IBM Press
Full text available: pdf(4.21 MB) Additional Information: full citation, abstract, references, index terms

Understanding distributed applications is a tedious and difficult task. Visualizations based on process-time diagrams are often used to o btain a better understanding of the execution of the application. The visualization tool we use is Poet, an event tracer developed at the University of Waterloo. However, these diagrams are often very complex and do not provide the user with the desired overview of the application. In our experience, such tools display repeated occurrences of no n-trivial commun ...

14 Aggregate nearest neighbor queries in spatial databases

Dimitris Papadias, Yufei Tao, Kyriakos Mouratidis, Chun Kit Hui
June 2005 **ACM Transactions on Database Systems (TODS)**, Volume 30 Issue 2
Publisher: ACM Press
Full text available: pdf(3.84 MB) Additional Information: full citation, abstract, references, index terms

Given two spatial datasets *P* (e.g., facilities) and *Q* (queries), an *aggregate nearest*

neighbor (ANN) query retrieves the point(s) of *P* with the smallest aggregate distance(s) to points in *Q*. Assuming, for example, *n* users at locations *q*1,...*qn*, an ANN query outputs the facility $p \in P$ that minimizes the *sum* of distances |*pq*i| for $1 \leqslant i \leqslant$...

Keywords: Spatial database, aggregation, nearest neighbor queries

15 Optimism and consistency in partitioned distributed database systems

Susan B. Davidson
September 1984 **ACM Transactions on Database Systems (TODS)**, Volume 9 Issue 3
Publisher: ACM Press

Full text available: pdf(1.88 MB) Additional Information: full citation, abstract, references, citings, index terms, review

A protocol for transaction processing during partition failures is presented which guarantees mutual consistency between copies of data-items after repair is completed. The protocol is "optimistic" in that transactions are processed without restrictions during failure; conflicts are then detected at repair time using a precedence graph, and are resolved by backing out transactions according to some backout strategy. The resulting database state ...

16 Management of disk space with REBATE

Shahram Ghandeharizadeh, Douglas J. Ierardi
November 1994 **Proceedings of the third international conference on Information and knowledge management**
Publisher: ACM Press

Full text available: pdf(981.06 KB) Additional Information: full citation, abstract, references, citings, index terms

The past decade has witnessed a proliferation of respositories whose workload consists of queries that retrieve information. These repositories provide on-line access to vast amount of data and serve as an integral component of many application domains (e.g., library information systems, scientific applications, entertainment industry). Their storage subsystem is expected to be hierarchical consisting of memory, disk drives, and one or more tertiary storage devices. The database resides per ...

17 An incremental access method for ViewCache: concept, algorithms, and cost analysis

Nicholas Roussopoulos
September 1991 **ACM Transactions on Database Systems (TODS)**, Volume 16 Issue 3
Publisher: ACM Press

Full text available: pdf(1.71 MB) Additional Information: full citation, abstract, references, citings, index terms, review

A ViewCache is a stored collection of pointers pointing to records of underlying relations needed to materialize a view. This paper presents an Incremental Access Method (IAM) that amortizes the maintenance cost of ViewCaches over a long time period or indefinitely. Amortization is based on deferred and other update propagation strategies. A deferred update strategy allows a ViewCache to remain outdated until a query needs to selectively or ...

Keywords: terms

18 Organization of clustered files for consecutive retrieval

J S. Deogun, V V. Raghavan, T K.W. Tsou

December 1984 **ACM Transactions on Database Systems (TODS)**, Volume 9 Issue 4
Publisher: ACM Press

Full text available: pdf(1.79 MB) Additional Information: full citation, abstract, references, index terms, review

This paper studies the problem of storing single-level and multilevel clustered files. Necessary and sufficient conditions for a single-level clustered file to have the consecutive retrieval property (CRP) are developed. A linear time algorithm to test the CRP for a given clustered file and to identify the proper arrangement of objects, if CRP exists, is presented. For the single-level clustered files that do not have CRP, it is shown that the problem of identifying a storage organization w ...

19 Object-based and image-based object representations

Hanan Samet
June 2004 **ACM Computing Surveys (CSUR)**, Volume 36 Issue 2
Publisher: ACM Press

Full text available: pdf(1.05 MB) Additional Information: full citation, abstract, references, index terms

An overview is presented of object-based and image-based representations of objects by their interiors. The representations are distinguished by the manner in which they can be used to answer two fundamental queries in database applications: (1) Feature query: given an object, determine its constituent cells (i.e., their locations in space). (2) Location query: given a cell (i.e., a location in space), determine the identity of the object (or objects) of which it is a member as well as the re ...

Keywords: Access methods, R-trees, feature query, geographic information systems (GIS), image space, location query, object space, octrees, pyramids, quadtrees, space-filling curves, spatial databases

20 Index-driven similarity search in metric spaces

Gisli R. Hjaltason, Hanan Samet
December 2003 **ACM Transactions on Database Systems (TODS)**, Volume 28 Issue 4
Publisher: ACM Press

Full text available: pdf(650.64 KB) Additional Information: full citation, abstract, references, citings, index terms

Similarity search is a very important operation in multimedia databases and other database applications involving complex objects, and involves finding objects in a data set S similar to a query object q, based on some similarity measure. In this article, we focus on methods for similarity search that make the general assumption that similarity is represented with a distance metric d. Existing methods for handling similarity search in this setting typically fall into one of ...

Keywords: Hiearchical metric data structures, distance-based indexing, nearest neighbor queries, range queries, ranking, similarity searching

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2006 ACM, Inc.
Terms of Usage Privacy Policy Code of Ethics Contact Us

Useful downloads: Adobe Acrobat QuickTime Windows Media Player Real Player

Search Session History

Fri, 13 Jan 2006, 10:44:31 AM EST

Edit an existing query or compose a new query in the Search Query Display.

Select a search number (#) to:

- Add a query to the Search Query Display
- Combine search queries using AND, OR, or NOT
- Delete a search
- Run a search

Search Query Display

Run Search Reset

Recent Search Queries

#1 ((partition database)<in>metadata)

#2 ((partition database)<in>metadata)

#3 ((partition database)<in>metadata) and file and context

#4 ((((partition database)<in>metadata) and file and context) <AND>(((partition database)) and file and context<in>metadata))

#5 (file and context and partition and database <IN>metadata)

#6 (((file and context and partition and database <in>metadata)) <AND>(location row and tuple<in>metadata))

#7 (partition database and value and lowest<IN>metadata)

#8 (partitioned database and location and value and file and context<IN>metadata)

Clear Session History

© Copyright 2005 IEEE –

Indexed by Inspec